Likes are not Likes
A Crowdworking Platform Analysis

Dennis Tatang[1], Philip Kreißel[2], Michael Sehring[1], Florian Quinkert[1], Martin Degeling[1], Thorsten Holz[1]
[1]Ruhr University Bochum, Germany
[2]Universität Bamberg, Germany

Abstract

Many Internet users rely on Online Social Networks (OSNs) in their daily lives to read news, find local restaurant recommendations, or learn about products. But lately OSNs have come under scrutiny for allowing misinformation to spread. Compared to the widely discussed problems of fake news and bots, *inauthentic likes by real users* have not seen much attention in the research community. They have their origin in real accounts that offer likes in exchange for micro-payments through crowd-working platforms. Our analysis of these fake likes from real users is based on almost 90,000 manipulation campaigns managed by a crowd-working company in Germany containing a diverse set ranging from politics to products and services. An additional study on ten users of these crowd-working platforms shows that they merely earn some pocket money and, although they do not like things they politically oppose, do not generally feel responsible for manipulations of recommendations. Our analysis shows that likes should not be trusted and the detection of fake accounts is not sufficient to prevent fake likes in general.

1 Introduction

Facebook, Twitter, and Instagram are among the largest online social networking sites. They are used by many Internet users on a daily basis and are an important medium for communication and information exchange. The largest social network Facebook, for example, reported 2.7 billion monthly active users during the second quarter of 2020 (Statista 2020) and the use of other social networks continues to grow. Social networks allow their users to interact with each other. Posts can be created and shared, and other users comment and like them. However, there is a growing concern that these interactions are not always genuine responses. Fake news, fake followers, and fake likes have become a real world problem in recent years (Xu et al. 2015; Stringhini et al. 2013). In particular, fake likes become harmful when messages are perceived as being popular, salient, or credible as a result of inauthentic interactions. Thus, for example, incorrect news about the current Coronavirus pandemic could be considered trustworthy when recommended by a large number of users. Various works dealt already with the issues of fake likes and fake accounts (De Cristofaro et al. 2014; Lin, Xia, and Liu 2015; Farooqi et al. 2017; Bay and Fredheim 2019). While the majority of these studies focused on the detection of automated bots some also looked at false reactions of real so-called crowdworkers (Ross et al. 2010; Wang et al. 2012).

In our work, we extend the knowledge regarding crowdworkers and crowdworker platforms. There are several providers that offer genuine likes or genuine followers as it is easy to identify some providers with a quick Web search (e.g., *https://buzzoid.com*, *https://stormlikes.com*, *https://www.social-viral.com*, and many more). Unfortunately, the exact number is not known but their low and competitive prices might suggest a large supply of similar products. Besides, we assume that the number of providers of these crowdworking platforms may grow, as bot detection by social networks is also improving. In this paper, we analyze one exemplary crowdworking platform that offers inauthentic likes for Facebook and other OSNs. We study this phenomenon from two perspectives: we report on like-campaigns purchased via the crowdworking platform and surveyed a small number clickworkers working for such platforms. We analyze 88,830 campaigns from a German crowdworking platform where real users earn money by liking content. Our analysis shows that representatives of almost all political parties in Germany, but also many different service providers, as well as products have social media appearances that were manipulated by crowdworkers of this platform. Through buying likes, these pages are likely perceived as more genuine and trustworthy (Chang, Yu, and Lu 2015). Additionally, we conduct a survey and interviews with crowdworkers to understand their reasoning for working at crowdworking platforms.

In summary, we make the following two main contributions:

- We study purchased likes and analyze all campaigns of a crowdworking platform for the first time.
- A small study on crowdworkers revealed that although it is not possible to earn more than a few Euros a month, they have been active for years without getting more than temporarily blocked.

2 Background

2.1 Online Social Network (OSN)

Online social networks (OSNs) or social network services (SNS) started out with a focus on maintain interpersonal relationships but have grown to information hubs that created ecosystems on their own. Major services include Facebook, Twitter, YouTube, Vkontakte, LinkedIn, and many others. Typical features of these platforms are the creation of a personal account with privacy settings to control what information (groups of) other users can access, a contact list, message exchange, as well as publishing and interacting with posts, images, or videos (Ellison, Steinfield, and Lampe 2007). At the time of writing, Facebook is the largest OSN (Statista 2020). Facebook also owns the video and photo-sharing application Instagram, as well as the messenger WhatsApp. Facebook states that more than three billion people around the world use their products (Facebook 2020a). The number of monthly active users on the Facebook platform is continuously growing since 2008 and reached 2.7 billion users in the second quarter of 2020 (Statista 2020).

Facebook Accounts To register a Facebook account users are required to enter their first name, last name, email address or mobile phone number, and a password. Users can voluntarily provide additional information and data, such as a profile pictures, a title picture, the place of residence, etc.

Facebook Pages Facebook offers the possibility to create pages for companies, brands, groups, or people of public interest. Users can like these pages, thus signaling they are interested in the respective content. It is possible to exchange private messages and to like and share posts from other people and pages. The number of likes of these pages are often seen as a sign of popularity and credibility (Facebook 2020c).

Facebook Likes On Facebook, user can mark different types of contributions with a *like*. In particular, we distinguish between three different types of likes:

1. *Page likes* are reactions to pages.
2. *Post likes*, on the other hand, are reactions that refer to contributions from pages or users (e.g., photos). The more reactions a post triggers, the higher is its potential reach (Facebook 2020e).
3. *Comment likes* are reactions to comments made on posts. The more reactions a comment has, the more visible this comment becomes under a posting (Facebook 2020d).

Facebook users can see the number of likes for pages, comments, and likes. Therefore many users might interpret likes as a sign of popularity and aspire themselves to acquire more likes, potentially leading to more traffic and ad revenue for Facebook.

2.2 Challenges in OSNs

Besides their advantages, OSNs also pose a number of challenges for the modern online society. On the one hand, the business model of OSNs is making money from user information, which may lead to privacy problems (Fire, Goldschmidt, and Elovici 2014). Recent headlines confirm that Facebook, e.g., is repeatedly involved in data-related incidents, where private data was used to target highly personalised ads to users based on their likes (Wong 2019; Winder 2019). Another recent scandal involved the creation of facial recognition technology based on scraped Facebook pictures (Hill 2020). On the other hand, OSNs are also be used to spread extremist or propagandistic viewpoints. Hate speech and fake news through social media are quick and easy to spread.

Fake accounts and also fake reactions can be used to increase the perceived popularity and credibility of these harmful contributions. OSN users are also exposed to opinion manipulation (Bradshaw and Howard 2019; Mihaylov, Georgiev, and Nakov 2015). As a result, a business model evolved in which likes and fake profiles can be purchased online. There are crowdsourcing systems that engage many users to solve a particular problem (Doan, Ramakrishnan, and Halevy 2011). In case of Facebook, there are so-called crowdworking platforms, allowing single users to earn a little extra money for distributing likes. These crowdworkers highlight the problem of SNS that a small number of users can have a high impact by interacting much more often than normal users. In addition to crowdworking platforms, so called *like farms* evolved, where fake accounts are used to raise specific pages or posts to achieve a higher general visibility. Previous work has analyzed factors that influence why real users distribute fake news (Talwar et al. 2019). Even research to crowdworking platforms exist (see Section 7), however, our data source is unique and offers extended insights into a known problem. Fake Accounts highlight a different problem of SNS, that a single human being can control multiple accounts on the network. Facebook publishes quarterly reports on their actions against fake accounts. These reports show that Facebook removes billions of accounts every year, most of them shortly after registration. Facebook estimates the prevalence of fake accounts at about five percent of monthly active users (Facebook 2020b).

3 Approach

In this section, we introduce the technical aspects of gathering data on fake like campaigns, and how we found participants for our qualitative survey. In detail, we describe three approaches. In Section 3.1, we explain our experiment, which motivated us to analyze one crowdworking platforms in depth. In Section 3.2, we present our methodology, how we collected data from a crowdworking platform and in Section 3.3, we report on how we contacted participants for our qualitative study.

3.1 Buying Fake Likes

To better understand the Fake-Like ecosystem, we decided to buy comment likes to examine the profiles that provide the purchased likes. In collaboration with a German news page, we created four comments under months-old posts to limit exposure of the comments to actual humans to a minimum. These comments were not related to the posts and contained cooking recipes. The analyzed fake like services of De Cristofaro et al. (De Cristofaro et al. 2014) could not be analyzed anymore because they no longer exist or cannot be used for comment likes.

After a short Internet research, we found several providers for fake likes on Facebook (e.g., fbpostlikes.com, boostlikes.com, fastlykke.com, famoid.com, followerpackages.com, fame-booster.de, likeservice24.de). Only four of these services were available and also offer purchasing likes for comments instead of original posts. Many accounts seem to be rather new and are only created to distribute likes. With these accounts one could suspect that they are bots. But there have also been some very genuine accounts. These accounts are considerably older and seem to be used by real users, because further interactions with other accounts are observed. For these accounts a simple bot detection or fake account detection is no longer sufficient. After further research we came across crowdworking platforms. Thus, this first insight led us to take a closer look at crowdworking platforms in the following.

3.2 Likes from Crowdworking Platform

Besides automated accounts and bots, "crowdworking platforms" are an important source of likes for services that want to increase their visibility in OSNs. In general, at a crowdworking platform, any Internet user can order a number of likes for an arbitrary website. Crowdworkers registered with that platform receive these orders, e.g. as a list, and like the websites accordingly to earn money. As part of our research, we reviewed various platforms and registered with a few of them to understand how they work and implement their platform. During our research, we discovered one particular supplier where a large number of campaigns of the last few years were easily accessible. This particular crowdworking platform lists current and previous campaigns, i.e., URLs to websites that are to be liked, on their platform allowing us to collect information about them. The structure of these URLs looked like this: *example-crowdworking-platform.de/*{*campaignID*}. The *campaignID* is simply incremented so that we could extract all websites (whole like list) that were to be liked since 2012 by sending requests to all valid campaignIDs. We knew that entries were going to start falling by 2012 because that was the year the crowdworking platform started operating. In particular, we requested campaignIDs between one (lowest possible campaignID) and 88,867 (highest campaignID at the time of performing the experiment). We determined the highest campaignID due to the simple fact that after this campaignID no more valid redirections to websites to be liked took place when we performed our analysis. Thus, in total, we identified 88,830 campaigns with 2,015 distinct targets, such as Facebook or YouTube. In 37 cases, the URLs were no longer available so that the number of identified campaigns differs from the highest campaignID.

This Germany-based platform uses the Facebook API to confirm that their crowdworkers actually liked the websites, which violates the terms of services for the API. Facebook claims to avoid misuse of the API by reviewing the particular use case before approving it. The example crowdworking platform circumvented this by submitting an application with a legitimate use case. Subsequently, they use the approved Facebook API to verify their crowdworkers' progress. This shows that malign actors can circumvent Facebook's API verification process with ease. As part of our work, we contacted Facebook and informed them that crowdworking platforms are misusing their API. Facebook responded that they do not approve selling fake reactions the business of selling likes. In the specific case of the German crowdworking platform, they reviewed the functionality of the app it used and then took it down, thus making it unavailable for use.

3.3 Recruitment of Fake Like Workers

To better understand how the economy around inauthentic likes works, we tried to interview and survey crowdworkers from multiple plattforms. We created two campaigns on different platforms and bought likes on Facebook and Twitter. We monitored the posts for the incoming likes and contacted the users through the respective channel to ask for their voluntary participation in a survey. This way we could be sure that they really are crowdworkers. Besides the ethical difficulties, it was also technically challenging to contact the accounts. Twitter and Facebook have strict rules for contacting other users. For example, on Facebook, messages by users that are not on one's friend list are often filtered. Twitter users have to allow messages from other users explicitly. In both cases, the platforms limit the number of messages we could send to recruit participants and sometimes blocked our accounts. Thus we could only send a limited number of messages per day. In some cases, we were then blocked for several hours or days.

Our participants could opt-in to a lottery of a 25 Euro gift card for completing the survey. In a pre-study, we determined that it took about 15 minutes to participate and complete the survey. We contacted 219 users in a total of which 10 (partly) completed our survey. The results below also include answers from two additional participants that had volunteered to participate in an in-depth interview. The low response rate might be a result of the platforms blocking or limiting notifications about our messages, but we also suspect that crowdworkers are aware that their work violate community rules of the platforms and is ethically questionable.

4 Results

Next, we present the results of our analysis of campaigns on the fake-like crowdworking platform and describe answers to our survey of crowdworkers.

4.1 Crowdworking Platform

Overview The vast majority of campaigns refers to Facebook (71,892 campaigns, 81 % of all campaigns). To a lesser extent, we found campaigns for other social networks (e.g., YouTube (4,833, 5.4 %) or Instagram (4,337, 4.9 %)) and other websites, like for example a German dating platform.

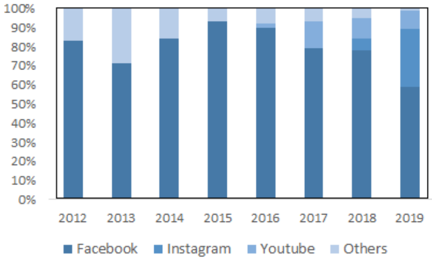

Figure 1: Share of found campaigns per platform per year.

When analyzing the campaigns over time, we see a shift between platforms of purchased Likes from Facebook to other social networks, such as Instagram. Figure 1 illustrates the share of each platform per year.

Facebook Campaigns In the following, we focus on the Facebook campaigns because the majority of the campaigns are still focused on that platform and it has the most users. In particular, we differentiate between likes for apps (25 campaigns), groups (60), photos (11,695), videos (1,831), posts (11,879), and pages/profiles (45,241).

Among the 25 identified apps were some well-known ones, such as Angry Birds or Candy Crush. We assume that supporters of these apps bought the likes or they were test campaigns of the crowdworking platform because it is rather unlikely that the developers of these apps used a crowdworking platform to buy likes. Additionally, we identified 60 distinct groups and 11,695 photos for which campaigns were started. The photos belong to 1,263 distinct Facebook pages with an average of nine photos per Facebook page and a median of one photo per page, indicating that the majority of these Facebook pages had manipulated the like-metrics of a low number of photos. A closer examination revealed that 982 of the Facebook pages are connected to less than five photos. In contrast, the high average is caused by a small portion of Facebook pages containing a comparatively high number of liked photos.

We found similar results for videos (1,831 videos for 497 distinct Facebook pages with an average of four videos, a median of one video, a maximum number of 172 videos for one Facebook page, and 434 Facebook pages with less than five videos) and posts (11,879 posts for 1,720 distinct Facebook pages, average of seven posts, median of one, a maximum of 393, and 1,327 Facebook pages with less than five posts).

The largest group of 45,241 campaigns is targeted at pages/ profiles, i.e., it is most common to buy likes for a page/profile. Additionally, we found pages/profiles which did not only buy likes for their pages/profiles, but also for photos and/or posts to make them appear more popular.

We identified Facebook pages that bought likes via the crowdworking platform covering a variety topics including political organisations. The list of campaigns contains Facebook pages of local chapters of all parties currently represented in the German federal parliament, see Figure 2.

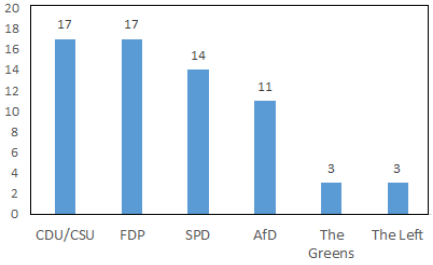

Figure 2: Number of social media appearances included in our campaign list, per party.

Table 1: Demographic information on our survey participants. (Note: It includes only selected answers.)

Gender		Age in years		Educational level		Income per month in Euro	
Male	5	>65	1	High School	2	2000 - 3000	1
Female	3	45 - 54	2	Secondary School	4	1000 - 2000	3
-	3	35 - 44	2	No school diploma	1	<500	2
		25 - 34	2	-	4	no own income	1
		-	4			-	4

Summary Our analysis of 88,830 manipulation campaigns shows that crowdworking platforms can offers fake likes on all social networks and that a variety of customers, including political parties, make use of these inauthentic user engagements. This indicates that crowdworking platforms for fake likes play a role in the campaigns to manipulate democratic processes. In the past social media companies have focused on limiting the influence of external involvement, e.g., by prohibiting organizations from outside of a country to purchase ads. Crowdworking platforms like the one we studied can provide an indirect way to manipulate how influential certain topics become within public debates. By increasing the number of likes certain topics might seem more relevant without disclosing, e.g., external funding to the platform.

4.2 Fake Like Workers

To learn more about the motivation of users to participate in these fake like schemes, we distributed a survey to individual that provided likes in the campaigns we initiated on two platforms.

Survey Overview The survey consisted of 33 questions on multiple topics. First, we wanted to know how participants conduct their work, where they find campaigns, and how many accounts they used. The second section contained questions on campaigns they had worked for and what types of posts and pages they did like in the past or would not like. In a third section, we asked participants about their opinions on fake likes and the online ratings in general. Last we asked them to answer some demographic questions.

Results Table 1 summarizes our demographic results of the survey. All participants where located in Western Europe. Despite the small number of responses, the demographics of our participants were diverse. Participants that answered the question about their gender (n=8) considered themselves equally male or female. Of the seven that disclosed their age, three where older than 45 years and from all educational backgrounds (no degree to high school diploma). The participants also disclosed their gross income ranging from less than 500 Euros per month (3), 1.000–2.000 (3) and 2.000–3.000 (1).

Regardless of their income, six participants said their motivation to create fake likes was financial, although the maximum amount any participant would earn was 21 to 50 Euros per month.

> It's an easy way to earn [...] some extra money. And it's also fun. I think it's interesting to like all those pages.

Seven participants disclosed that they are earning 20 Euros or less per month. Two participants answered they were doing it to discover new things or to kill time. The low income from this work corresponds with the amount of time participants spent clicking for campaigns. Seven out of eight that answered the question reported that they spent between 11 and 30 minutes per day in which four click up to 50, two up to 100, and two more than one hundred campaigns. Even in the best case scenario of 11 minutes per day, this would mean that these workers earn less than the minimum wage on this job which is consistent with reports for similar platforms like Mechanical Turk (DeSoto 2016). Nevertheless, three out of the six participants that reported that they have been doing this work for more than three years.

Participants in interviews reported that they, on average, get about 2 Euro cents per click. However, income possibilities are not only limited by the number of campaigns available but also by the platforms themselves. The platform analyzed above limits the number of clicks to a maximum of 40 per day, presumably to ensure that users do not get banned or marked as bots. Still, four of nine participants that responded to the survey reported at least one of their accounts had been blocked because of their activity.

> I think I clicked too fast. The [crowdworking] guidelines say that you should stay a minute on the page you clicked. In the beginning I was euphoric and clicked one after the other. [...] I was blocked afterwards. For a day or two, so [my acccount] was not really closed.

Regardless, the majority of participants use their private account to conduct the fake likes. Only one participant reported that she uses multiple separate accounts for the clickworking job.

We also asked participants about the types of pages and posts of campaigns. We presented them with a category schema similar to Facebooks page categorization[1] and presented participants with a five-points scale from never (1) to always liking those pages (5). While the averages for most of the categories where similarly high (avg $\geq$ 4.4, n=8) (Businesses, Media, Public Figures, NGOs) respondents reported to less often like pages or posts of political parties (average = 3.2, n=8). Correspondingly, five out of seven participants reported to have refrained from liking specific content in the past. While one participant responded to not really look at the content, three specifically mentioned to avoid political parties and posts with right-wing or racist content. Two participants also mentioned that they do not really look at the content, but try to avoid pages with content for kids or pornography.

Given the public debate about manipulation online as well as the personal experience with producing fake likes, we also asked the participants about their experience with online reviews and likes and how much they trust assessments made by other web users. To our surprise, four out of seven respondents to the question replied their view on reviews and likes has not changed due to their work. While two reported they now have a negative view, one response even claimed that their view is now much more positive. We than again asked questions regarding the trustworthiness of online reviews and likes in specific categories on a five points scale. From the answers of six participants general news was rated highest (avg = 3.8) while advertisements received the lowest score (2.8), in between where "political news", "product recommendations", and "other".

We then wanted to learn more about the participants perception of their own responsibility in the context of distorting the online reviews by asking "What do you think about the statement that purchased Likes falsify the rating system on online platforms?". Four out of six participants chose the neutral answer here. In one of the interviews, one person expanded on the topic and claimed that there is no truth online, anyway:

> I've become sceptic. In the past I was a fan of [a major news show], but since I started reading more about things I had to learn that nearly all news are fake.

Summary The results of our survey are limited due to the low response rate, although we received answers from a variety of participants with respect to their demographics. The answers show that none of the participants produces fake likes for a living, although all see the additional income as their primary motivation. The income is limited due to the social media platforms trying to prevent this type of use but seem to consistently fail as no participants reported to have been continuously blocked, but instead have been doing this for years.

While the public debate often centers around manipulation for political purposes, it seems that the majority of campaigns are focused on businesses and products. Well aware of the problem with political advertising, the majority of participants reported that they do not click the like button for political parties or representatives they do not approve off. Although the majority has an unfavorable view on review systems, they do not necessarily feel responsible for misleading others. Instead, they emphasize the responsibility of other users platforms showing a low level of online trust (Grabner-Kräuter and Bitter 2015).

4.3 Example Case Study: Garden Furniture Shop

Among the campaigns of the crowdworking platform, we found a page of a local German shop for garden furniture with over a million likes. We assume that this is highly unusual for such a small business. Even though the campaign

[1] see https://www.facebook.com/pages/category/

in the crowdworking platform only included likes for the page itself, we also found that the posts of this page received extremely high numbers of likes (up to many ten thousands). In this example, when we look at the accounts that the likes assigned to, it was quite obvious that the majority of these accounts did not have German names. Thus, we believe that most likes on this page are bought and thus forged. This example underlines that it is possible to operate a page on Facebook, where the majority of the interactions with this page are not genuine.

5 Discussion

Lessons Learned Purchased likes from like farms or crowdworking platforms are a known problem since users of online social networks can be deceived by purchased likes. Our results underline that not only bots and fake accounts are responsible for fake likes, but also genuine users, who distribute fake likes as crowdworkers on different platforms for a little amount of money. Our tests show that it is very easy for third parties to purchase likes on platforms and direct crowdworkers to specific posts - even without being involved in the original post and without leaving a trace on the platform. It is almost impossible to distinguish crowdworkers from normal users because they are genuine accounts. The fact that all major political parties make use of fake like campaigns shows that this is a widespread problem that could become increasingly problematic. Thus, especially in times of election campaigns, the number of likes should always be critically assessed.

Proposed Solutions SNSs such as Facebook should invest more in the authenticity of their user base. It is far too easy to buy likes from inauthentic-looking accounts and the number of user registrations does not seem to be under control. Even if this problem was solved, authentic users could still sell their activities on the network to crowdworking platforms by liking multiple pages per day. This risks giving users a wrong sense of popularity for potentially harmful and political messages. Therefore, Facebook should hide the like count from public view until these issues are under control. Instagram already begun with first experiments with hidden like counts (Yurieff 2019).

Further Research Future research could try to extend our work (e.g, by finding other crowdworking pages with similar ID-Systems). But our research might suggest, that it may be unlikely to find many pages with such vulnerabilities. However from our experience in this paper we can suggest a different approach to reach similar results than our own: Many crowdworker-profiles had very limited privacy settings, thus making it possible to compare their page likes. Analysing a large number of accounts bought from a single company might reveal patterns of frequently liked pages. This method is even simpler on Twitter, where timeline and accounts followed can be retrieved via an API. It is possible to identify an initial group of crowdworkers by buying follows, retweets or likes from a crowd-working plattform.

Threats To Validity Our findings regarding the crowdworking platform are based on a single German crowdworking platform. A general statement is therefore only possible to a limited extent. However, so far there was no possibility to collect such information about such platforms at all, so our first insight is interesting. We also offer ideas to expand on our approach. We have also only been able to survey a limited number of participants in our qualitative survey. Nevertheless, we dive into this topic and provide some interesting and useful insights complementary to other works. We hope that in future work these insights can be evaluated in more detail (e.g., to understand how these crowdworking platforms are able to influence online discourse).

6 Ethical Considerations

For this work, we collected and analyzed information worthy of protection. Our research institution does not require approval or provide an institutional review board (IRB) for this type of studies. Instead, we consulted our universities data protection officer and followed common guidelines of the research community to protect those whose data was collected (Bailey et al. 2012). Particularly the collected data of the German crowdworking company has been stored securely. We also contacted the operator of the crowdworking company and he gave us the permission to publish the information. Nevertheless, we will not mention the name of the company in this paper. After completing the data collection of the crowdworking platform, we also informed Facebook about our research. They responded positively and suspended the exploited service. We hope that our research will help to raise awareness for the problem of paid likes in order to conquer it. When contacting clickworkers for interviews, we made sure to protect their privacy. The clickworkers received a link to the survey and a financial incentive, there was no obligation to take part in the study. The data was minimized and state-of-the-art mechanisms were used to protect the data in transmission and rest. All gathered data was collected for scientific purposes only and will not publicly available (upon request parts of our data might be shared).

7 Related Work

Various works explored social networks, fake likes, fake accounts and also crowdworking platforms. In the following we provide a brief insight into these research areas.

An early work on crowdworkers is conducted by Ross et al. from 2010 (Ross et al. 2010) who analyze worker demographics from the Mechanical Turk platform. A closely related work is by Wang et al. from 2012 (Wang et al. 2012), the authors examined in particular two crowdworking platforms in China. Their conclusion is that crowdsourcing systems are a global problem. Our analysis confirms that crowdsourcing systems are still widely used almost a decade later. In contrast to their work, our analysis is based on a non-public data set. Another work, which also deals with crowdworking platforms, is by Rinta-Kahila and Soliman from 2017 (Rinta-Kahila and Soliman 2017), who examined the concept of crowdturfing respectively crowdworking and conducted a literature review. We contribute to this knowledge with our study and hope to draw attention to the topic in society in order to counteract crowdturfing. A more recent work from 2018 that also looks at services to manipulate social networks has been done by DeKoven et al. (DeKoven et al. 2018). The authors took a closer look at the Instagram platform and analyzed 5 manipulation services. In 2018, there was also another study on YouTube (Hussain et al. 2018). We believe that the many works in this area show that this is a relevant problem. Our analysis fits well into the series of these works and especially the aspect of our survey with crowdworkers offers new insights in this research area.

Many other papers also deal with the analysis of different general aspects of OSNs. As early as 2010, spam campaigns on Facebook were analyzed for the first time by Gao et al. (Gao et al. 2010). In 2011 Liu et al. (Liu et al. 2011) analyzed Facebook privacy settings. For that, the authors conducted a survey with 200 participants.Stutzman et al. (Stutzman, Gross, and Acquisti 2013) analyzed a set of more than 5,000 profiles at CMU between 2005 and 2011, documenting how information disclosure of users evolved over time. Garcia et al. (Garcia-Gavilanes, Quercia, and Jaimes 2013) examined cultural differences in the use of social networks in 2013. To this end, the authors have focused on Twitter. In the work by De Cristofaro et al. (De Cristofaro et al. 2014) from 2014, the authors bought fake likes for honeypot Facebook pages and then analyzed them. This work is comparable to our purchased comments likes, but we update this work and extend it with previously unknown information not only to like farming but also to crowdworking platforms. In 2015, Aggarwal et. al. (Aggarwal and Kumaraguru 2015) analyzed markets for semi-automated followers on Twitter. While real accounts can participate in these markets, the follow-process itself is automated via Twitter API. This is different from the crowdworking plattforms we analyzed: Here likes are distributed by manual workers. In 2015, Lin et al. (Lin, Xia, and Liu 2015) presented a detailed study of the like button of Facebook. The authors showed that fake likes can be easily created. They analyzed the number of likes of more than 9,000 websites and could identify fake like buyers with abnormal like number increase. In 2017, Farooqi et al. (Farooqi et al. 2017) analyzed collusion networks on Facebook on a large-scale dataset gathered via honeypot Facebook accounts. The authors discussed so far overlooked reputation manipulation services on Facebook. A recent work, which also analyzes Facebook as OSN, is by Andreou et al. (Andreou et al. 2019). The authors present a study on the Facebook advertising ecosystem. We consider a completely different aspect. The work of Sen et al. (Sen et al. 2018) from 2018 considers fake likes at Instagram. They believe that false engagement needs to be looked at more closely. There may be ordinary users who distribute both fake and genuine likes. Exactly in this area we researched with our work. Bay et. al. (Bay and Fredheim 2019) which work for the NATO Stratcom bought more than 54,000 fake social media reactions on multiple platforms. They found that SNSs failed to remove the fake reactions as four in five fake reactions were still online after four weeks.

In our work, we focus on crowdworking platforms, their campaigns and their workers, especially in the analysis of fake likes on OSNs and the reasons for being a crowdworker. Thereby, we contribute additional information to the knowledge of how crowdworking systems work in general.

8 Conclusion

In summary, we showed by the analysis of a crowdworking platform that fake likes belong not only to fake accounts but also to crowdworkers and that all sectors of pages contain purchased fake likes via this platform. One way to avoid the problems with fake likes in the future, OSNs should remove the like count. Facebook switched off the like information on Instagram for individual accounts in first pilot projects and plans to transfer it to Facebook under certain circumstances (Constine 2019).

References

Aggarwal, A., and Kumaraguru, P. 2015. What they do in shadows: Twitter underground follower market. In *2015 13th Annual Conference on Privacy, Security and Trust (PST)*.

Andreou, A.; Silva, M.; Benevenuto, F.; Goga, O.; Loiseau, P.; and Mislove, A. 2019. Measuring the Facebook Advertising Ecosystem. In *NDSS 2019 - Proceedings of the Network and Distributed System Security Symposium*.

Bailey, M.; Dittrich, D.; Kenneally, E.; and Maughan, D. 2012. The menlo report. *IEEE Security & Privacy*.

Bay, S., and Fredheim, R. 2019. How Social Media Companies are Failing to Combat Inauthentic Behaviour Online. https://www.stopfake.org/en/how-social-media-companies-are-failing-to-combat-inauthentic-behaviour-online/. Accessed: 2020-01-20.

Bradshaw, S., and Howard, P. N. 2019. The global disinformation disorder: 2019 global inventory of organised social media manipulation. *Project on Computational Propaganda*.

Chang, Y.-T.; Yu, H.; and Lu, H.-P. 2015. Persuasive messages, popularity cohesion, and message diffusion in social media marketing.

Constine, J. 2019. Now Facebook says it may remove Like counts. https://techcrunch.com/2019/09/02/facebook-hidden-likes/amp/. Accessed: 2020-07-30.

De Cristofaro, E.; Friedman, A.; Jourjon, G.; Kaafar, M. A.; and Shafiq, M. Z. 2014. Paying for likes?: Understanding facebook like fraud using honeypots. In *Proceedings of the 2014 Conference on Internet Measurement Conference*.

DeKoven, L. F.; Pottinger, T.; Savage, S.; Voelker, G. M.; and Leontiadis, N. 2018. Following their footsteps: Characterizing account automation abuse and defenses. In *Proceedings of the Internet Measurement Conference 2018*.

DeSoto, K. A. 2016. Under the Hood of Mechanical Turk. *APS Observer*.

Doan, A.; Ramakrishnan, R.; and Halevy, A. Y. 2011. Crowdsourcing systems on the world-wide web. *Communications of the ACM*.

Ellison, N. B.; Steinfield, C.; and Lampe, C. 2007. The Benefits of Facebook "Friends:" Social Capital and College

Students' Use of Online Social Network Sites. *Journal of Computer-Mediated Communication*.

Facebook. 2020a. facebook Newsroom - Company Info. https://newsroom.fb.com/company-info/. Accessed: 2020-07-30.

Facebook. 2020b. Facebook Transparency - Fake Accounts. https://transparency.facebook.com/community-standards-enforcement\#fake-accounts. Accessed: 2020-07-30.

Facebook. 2020c. Organic Reach on Facebook: Your Questions Answered. https://www.facebook.com/business/news/Organic-Reach-on-Facebook. Accessed: 2020-07-30.

Facebook. 2020d. What does Most Relevant mean on a Facebook Page post? https://www.facebook.com/help/539680519386145. Accessed: 2020-07-30.

Facebook. 2020e. What influences the order of posts in my Facebook News Feed? https://www.facebook.com/help/520348825116417. Accessed: 2020-07-30.

Farooqi, S.; Zaffar, F.; Leontiadis, N.; and Shafiq, Z. 2017. Measuring and mitigating oauth access token abuse by collusion networks. In *Proceedings of the 2017 Internet Measurement Conference*.

Fire, M.; Goldschmidt, R.; and Elovici, Y. 2014. Online social networks: threats and solutions. *IEEE Communications Surveys & Tutorials*.

Gao, H.; Hu, J.; Wilson, C.; Li, Z.; Chen, Y.; and Zhao, B. Y. 2010. Detecting and characterizing social spam campaigns. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*.

Garcia-Gavilanes, R.; Quercia, D.; and Jaimes, A. 2013. Cultural dimensions in twitter: Time, individualism and power. In *Seventh International AAAI Conference on Weblogs and Social Media*.

Grabner-Kräuter, S., and Bitter, S. 2015. Trust in online social networks: A multifaceted perspective. *Forum for Social Economics*.

Hill, K. 2020. The Secretive Company That Might End Privacy as We Know It. https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html. Accessed: 2020-01-20.

Hussain, M. N.; Tokdemir, S.; Agarwal, N.; and Al-Khateeb, S. 2018. Analyzing disinformation and crowd manipulation tactics on youtube. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*.

Lin, X.; Xia, M.; and Liu, X. 2015. Does" like" really mean like? a study of the facebook fake like phenomenon and an efficient countermeasure. *arXiv preprint arXiv:1503.05414*.

Liu, Y.; Gummadi, K. P.; Krishnamurthy, B.; and Mislove, A. 2011. Analyzing facebook privacy settings: user expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*.

Mihaylov, T.; Georgiev, G.; and Nakov, P. 2015. Finding opinion manipulation trolls in news community forums. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*.

Rinta-Kahila, T., and Soliman, W. 2017. Understanding crowdturfing: the different ethical logics behind the clandestine industry of deception.

Ross, J.; Irani, L.; Silberman, M. S.; Zaldivar, A.; and Tomlinson, B. 2010. Who are the crowdworkers? shifting demographics in mechanical turk. In *CHI'10 extended abstracts on Human factors in computing systems*.

Sen, I.; Aggarwal, A.; Mian, S.; Singh, S.; Kumaraguru, P.; and Datta, A. 2018. Worth Its Weight in Likes: Towards Detecting Fake Likes on Instagram. In *Proceedings of the 10th ACM Conference on Web Science*.

Statista. 2020. Number of monthly active Facebook users worldwide as of 2nd quarter 2020 (in millions). https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/. Accessed: 2020-09-02.

Stringhini, G.; Wang, G.; Egele, M.; Kruegel, C.; Vigna, G.; Zheng, H.; and Zhao, B. Y. 2013. Follow the green: growth and dynamics in twitter follower markets. In *Proceedings of the 2013 conference on Internet measurement conference*.

Stutzman, F.; Gross, R.; and Acquisti, A. 2013. Silent Listeners: The Evolution of Privacy and Disclosure on Facebook. *Journal of Privacy and Confidentiality*.

Talwar, S.; Dhir, A.; Kaur, P.; Zafar, N.; and Alrasheedy, M. 2019. Why do people share fake news? Associations between the dark side of social media use and fake news sharing behavior. *Journal of Retailing and Consumer Services*.

Wang, G.; Wilson, C.; Zhao, X.; Zhu, Y.; Mohanlal, M.; Zheng, H.; and Zhao, B. Y. 2012. Serf and turf: crowdturfing for fun and profit. In *Proceedings of the 21st international conference on World Wide Web*.

Winder, D. 2019. Unsecured Facebook Databases Leak Data Of 419 Million Users. https://www.forbes.com/sites/daveywinder/2019/09/05/facebook-security-snafu-exposes-419-million-user-phone-numbers/\#29c23b4f1ab7. Accessed: 2020-07-30.

Wong, J. C. 2019. The Cambridge Analytica scandal changed the world - but it didn't change Facebook. https://www.theguardian.com/technology/2019/mar/17/the-cambridge-analytica-scandal-changed-the-world-but-it-didnt-change-facebook. Accessed: 2020-07-30.

Xu, H.; Liu, D.; Wang, H.; and Stavrou, A. 2015. E-commerce reputation manipulation: The emergence of reputation-escalation-as-a-service. In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee.

Yurieff, K. 2019. Instagram is now testing hiding likes worldwide. https://edition.cnn.com/2019/11/14/tech/instagram-hiding-likes-globally/index.html. Accessed: 2020-09-15.